

REMARKS

Applicant appreciates the time taken by the Examiner to review Applicant's present application. This application has been carefully reviewed in light of the Official Action mailed January 2, 2009. Applicant respectfully requests reconsideration and favorable action in this case.

Claims Status

Claims 1-28 and 59-72 were rejected. Claims 29-59 were withdrawn. Claims 1, 2, 6-8, 10-11, 16-17, 20, 23, 26, 60, 64 and 71 are amended herein. Claims 3-5, 9, 12-15, 19, 21-22, 24-25, 27, 62, 65-69 and 72 are cancelled. Support for the amendments may be found in the specification at least in paragraphs 63, 66-67, 70, 73 and 80. No new claims are added. Thus, claims 1-2, 6-8, 10-11, 16-18, 20, 23, 26, 28, 60-61, 63-64, and 70-71 remain pending.

Specification Objections

The specification stands objected to as failing to provide proper antecedent basis for the claimed subject matter. The claims are amended. Accordingly, Applicant submits the objection is moot.

Rejections under 35 U.S.C. § 103

Claims 1-3 and 6-28 were rejected as obvious over U.S. Patent No. 6,498,786 ("Kirkby") in view of U.S. Publication No. 2003/0126202 ("Watt"). Applicant traverses the rejection.

Claim 1, as amended, recites:

A system for allocating resources amongst a plurality of applications running in a distributed, multiprocessor computing environment, the system comprising:
a plurality of server computers in a server pool, wherein each server comprises;
a distributed server pool director for organizing and maintaining the plurality of servers in the server pool;
a monitoring module at each server computer in the server pool for detecting demands for one or more resources located on the server computer and exchanging information amongst the plurality of server computers regarding demands for the one or more resources at the plurality of servers;
a distributed policy engine for specifying a policy for allocation of resources of the plurality of server computers amongst the plurality of applications having access to the resources, wherein the distributed policy engine bases the policy on the demands for the one or more resources, monitors the system to determine when one or more conditions exist, and adapts the policy based on the one or more conditions, and wherein the policy

engine is distributed on each server computer in the plurality of server computers; and
an enforcement module at each server computer for allocating the resources amongst the plurality of applications based on a decision made by the policy engine.

Thus, embodiments of a system for allocating resources amongst a plurality of applications may include a plurality of server computers in a server pool, a distributed pool server director for organizing and maintaining the set of server computers in the server pool, a monitoring module at each server computer in the server pool for detecting demands for one or more resources located in the server computer and exchanging information regarding demands for the one or more resources at the plurality of servers, a distributed policy engine for specifying a policy for allocation of resources of the plurality of server computers amongst the plurality of applications having access to the resources, and an enforcement module at each server computer for allocating the resources amongst the plurality of applications based on a decision made by the policy engine. The distributed policy engine can base the policy on the demands for the one or more resources, monitor the system to determine when one or more conditions exist, and adapt the policy based on the one or more conditions.

Kirkby describes a method of allocating bandwidth in a telecommunications network. As mentioned in the Reply to Office Action dated July 30, 2008 Kirkby describes a system which receives customer input and determines a Willingness-to-Pay value (WtP), such that allocation of resources is made based on a customer's willingness to pay for more resources. (See, Response to Office Action dated July 30, 2008, page 12, lines 20-23.) A network manager determines which resources a user will use when sending traffic to a selected destination node. (See, Kirkby, Col. 10, lines 10-12.)

Watt describes a system and method for dynamic server allocation and provisioning, in which server images are stored and made available for loading onto one or more servers in a data center, a load manager assigns one of the server images to each of the servers, and at least one server manager is associated with each of the servers capable of powering it on and off. Watt further describes a server monitor that receives periodic heartbeat and load measurement signals from each server in the data center and then reports the signals to the load manager, which allocates and provisions servers upon detecting failures (i.e. lack of heartbeat signals from a particular server in the data center). (See, Watt, paras. 12-13.)

In the rejection, the Examiner states that it would have been obvious to combine the server pool management taught by Watt with the policy-based resource management in distributed systems taught by Kirkby in order to streamline resource allocation and provisioning. Applicant respectfully submits that both the method of bandwidth allocation taught by Kirkby and the method of allocating servers taught by Watt teach centralized systems and that the combination would not teach or anticipate a distributed system for resource allocation.

To implement Kirkby's invention, it is necessary that a central network manager (such as network manager 50 in Figure 1) communicates with each network element manager. (See, Kirkby, Col. 9, line 66 – Col. 10, line 7.) In particular, using Kirkby's system, a method for controlling bandwidth resources may be either applied by a centralized network manager or delegated by the network manager to controllers that are distributed on the network. (See, Kirkby, Col. 3, lines 47-52.) In the case that the bandwidth allocation is handled by a network manager, Kirkby teaches that in some cases, centralized control by the network manager is desirable since the manager has time to poll each of the users for WtP values and to then allocate bandwidth accordingly. (See, Kirkby, Col. 11, lines 17-21.) In the case that the method is delegated to local network element managers, bandwidth is allocated to local controller and then the local network element manager allocates bandwidth to users. For example, as described by Kirkby, the overall negotiation of network traffic contracts may be coordinated by a corporate network manager (90) based on high level corporate policies. Each customer premises equipment manager (CPE) (60) allocates a share of bandwidth for all users in proportion to the WtP values set by the corporate policy. (See, Kirkby, Col. 11, line 34 – Col. 12, line 15.) This appears to be a multi-tiered approach to resource allocation. Thus, Kirkby teaches either a single centralized controller or a multi-tiered approach which ultimately has a central controller delegating to local controllers.

In contrast, embodiments disclosed by Applicant do not rely on a centralized network manager to allocate resources, but instead utilize a distributed system to allocate resources. The system includes a policy engine, a monitoring module, a server pool director and an enforcement module, each of which run on every server in the server pool. (See, specification, para. 70.) The policy engines distributed on the servers examine a number of factors (e.g., resource availability and resource demands by applications) and their interdependencies and then apply fine-grained policies for allocation of resources, such as making decisions about actions to be taken on a

request. (See, specification, paras. 66 and 73.) As a result, there is no need for a centralized controller to communicate with each server to allocate every resource.

Applicant respectfully submits that the shortcomings of Kirkby are not remedied by the teachings of Watt. Watt appears to teach that the components are not distributed across the plurality of servers. For example, Figure 2 depicts server pools 212a-212c as separate from load manager 206, server monitor 204 and network storage element 218. Furthermore, Watt's description of the components depicted in Figure 2 seem to teach a centralized system. Namely, the repository manager is responsible for provisioning and managing server images such that an administrator can install software one time and it can be shared by all servers, that each server manager is responsible for rebooting or powering off specified servers as directed by a load manager, that a server monitor monitors the health, load and response time of all servers and that a load manager is responsible for allocating servers and images. (See, Watt, paras. 42-49.) The description of these components and the lack of any language to the contrary appears to indicate that Watt is not teaching a distributed system. Additionally, Watt also describes a control console (not shown in Figure 2) that allows data center personnel to create server images, manage the software repository, manually allocate servers and images, etc. (See, Watt, para. 53.) Applicant submits that the level of manual control suggested by Watt is accomplished through a centralized controller. Thus, Applicant submits that Watt teaches using centralized control in a data center that may include some level of manual control.

In contrast, embodiments disclosed by Applicant use a distributed system and method for automatically allocating resources. A distributed server pool director organizes and maintains the set of servers in the server pool. (See, specification, para. 74.) Advantageously, after a user provides an initial policy or a set of priorities, a policy engine distributed on each server in the system is able to work out low level implementation details without further user direction and reallocate the resources accordingly. Thus, allocation of resources may be automatically controlled by components distributed across the servers.

For at least the foregoing reasons, Applicant respectfully submits that Kirkby's allocation of bandwidth and Watt's dynamic allocation of servers, alone or in combination, fail to teach or describe a distributed system for allocating resources. Accordingly, withdrawal of this rejection is requested.

Claims 4 and 5 were rejected as obvious over Kirkby in view of U.S. Patent No. 5,414,845 ("Behm"). Claims 4 and 5 are cancelled. Applicant submits that the rejection is now moot.

Claims 60-64, 66-67, 70 and 71 were rejected as obvious over U.S. Patent No. 6,766,348 ("Combs") in view of U.S. Publication No. 2003/0069972 ("Yoshimura") and Watt. The rejection is traversed.

Combs is generally concerned with load-balanced data exchange in distributed network-based resource allocation. Combs describes a resource allocation system having a number of resource allocation system agents (RASAs), each running as a separate process on a computer connected to the network. The RASAs communicate with each other to provide a load balancing mechanism that allocates resources to users on the basis of similar domain and greatest remaining capacity. Combs teaches that allocation on the basis of closest domain balances communication resources and allocation on the basis of greatest remaining capacity evenly distributes utilization of similar resources. (See, Combs, Col. 2, lines 38-65.)

In the rejection, the Examiner states that Combs discloses a method of receiving user input specifying priorities of the plurality of applications to resources of a plurality of servers, selecting an application based on the specified priorities of the plurality of applications and allocating to the application resources based on the policy for allocation of resources. Applicant respectfully submits that the methods described by Combs concern allocated resources but do not concern the decision of how to allocate the resources. Combs is concerned with load balancing, such that if there are ten modems and Combs is given seven modems and three are allocated to other users, Combs may use load balancing between the seven to allocate the resource. (See, Combs, Col. 8, lines 37-44.) However, Combs is only concerned with load balancing between the seven modems allocated to him and fails to say how the ten modems are allocated. Thus, if Combs is given a set of seven modems out of a pool of ten modems and the seven are operating at capacity while the other three are idle, Combs is silent as to how one or more of the three modems may be re-allocated to the set of seven modems.

In contrast, Applicant discloses methods for allocating resources using a distributed system. As an example, assuming Applicant had ten modems, a user may give input specifying the priority of a number of applications and embodiments of Applicant's system may allocate three modems to a first set of applications and seven modems to a second set of applications. Furthermore, embodiments disclosed by Applicant may, at some time, evaluate the state of the ten servers and, as a result, change the number of modems allocated to an application, power

off one or more modems to reduce the number of modems, etc. Thus, for example, a system of ten servers may have seven servers allocated to sales for the majority of the month but may reallocate some of the servers at various times of the month to facilitate billing, inventory or the like.

Applicant submits that Yoshimura fails to remedy the deficiencies of Combs in that the load balancing of resources, as taught by Yoshimura, fails to teach determining how to allocate resources. Yoshimura teaches using a load balancer to select a computer from a set of computers configured to execute a user request. Furthermore, when there are plural load balancers, a managing server controls a table to be matched between the load balancers. (See, Yoshimura, para. 11.)

In contrast, the allocation of resources by Applicant is generally performed by a policy engine that is distributed on a plurality of servers, wherein the policy engine receives user input specifying priorities of the plurality of applications to resources of a plurality of servers, selects an application based on the specified priorities of the plurality of applications and allocates to the application resources based on the policy for allocation of resources. Thus, embodiments disclosed by Applicant may use a policy engine distributed across a plurality of servers to allocate resources instead of having a single managing server load balance between the resources.

For at least the foregoing reasons, Applicant submits that Combs and Yoshimura, alone or in combination, fail to teach a method of a policy engine distributed across a plurality of servers receiving user input specifying priorities of the plurality of applications to resources of a plurality of servers, selecting an application based on the specified priorities of the plurality of applications and allocating to the application resources based on the policy for allocation of resources, as recited in amended claim 60. Accordingly, withdrawal of this rejection is requested.

Claims 65, 68 and 69 were rejected as obvious over Combs, Yoshimura and Watt in view of U.S. Publication No. 2005/0177755 ("Fung"). Claims 65, 68 and 69 are cancelled. Accordingly, Applicant submits that this rejection is now moot.

Conclusion

Applicant has now made an earnest attempt to place this case in condition for allowance. Other than as explicitly set forth above, this reply does not include an acquiescence to statements, assertions, assumptions, conclusions, or any combination thereof in the Office Action. For the foregoing reasons and for other reasons clearly apparent, Applicant respectfully requests full allowance of Claims 1-2, 6-8, 10-11, 16-18, 20, 23, 26, 28, 60-61, 63-64, and 70-71. The Examiner is invited to telephone the undersigned at the number listed below for prompt action in the event any issues remain.

The Director of the U.S. Patent and Trademark Office is hereby authorized to charge any fees or credit any overpayments to Deposit Account No. 50-3183 of Sprinkle IP Law Group.

Respectfully submitted,

Sprinkle IP Law Group
Attorneys for Applicant

Ari G. Akmal
Reg. No. 51,388

Date: *May 28, 2009*

1301 W. 25th Street, Suite 408
Austin, TX 78705
Tel. (512) 637-9220
Fax. (512) 371-9088